




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁵ : H04L 12/56, H04J 3/24	A1	(11) International Publication Number: WO 93/06675 (43) International Publication Date: 1 April 1993 (01.04.93)
(21) International Application Number: PCT/US92/07978 (22) International Filing Date: 25 September 1992 (25.09.92) (30) Priority data: 765,776 26 September 1991 (26.09.91) US (71) Applicant: COMMUNICATIONS SATELLITE CORPORATION [US/US]; 950 L'Enfant Plaza, S.W., Washington, DC 20024 (US). (72) Inventor: SHYY, Dong-Jye ; 19729 Crystal Rock Dr., Apt. 14, Germantown, MD 20874 (US). (74) Agent: FOURNIER, Kevin, J.; Sughrue, Mion, Zinn, Macpeak & Seas, 2100 Pennsylvania Ave., N.W., Washington, DC 20037-3202 (US).		(81) Designated States: CA, JP, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, SE). Published <i>With international search report.</i>

(54) Title: NONBLOCKING MULTICAST FAST PACKET/CIRCUIT SWITCHING NETWORKS



(57) Abstract

A self-routing nonblocking multicast switching network (Fig. 4, Fig. 5, Fig. 6, Fig. 7, or Fig. 8) routes input messages to destined addresses by examining a routing tag (Fig. 3b, Fig. 3c and Fig. 3d) combined in each message. The routing tag (Fig. 3b, Fig. 3c and Fig. 3d) has a plurality of sections (301, 302, 303, 304) each section corresponding to a level of a tree hierarchy (Fig. 3a) related to the outputs of the switching network. The network sorts the messages by examining only one section of each routing tag and routes and sorted messages to the destined addresses based on the bits contained in the routing tags.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	FI	Finland	MN	Mongolia
AU	Australia	FR	France	MR	Mauritania
BB	Barbados	GA	Gabon	MW	Malawi
BE	Belgium	GB	United Kingdom	NL	Netherlands
BF	Burkina Faso	GN	Guinea	NO	Norway
BG	Bulgaria	GR	Greece	NZ	New Zealand
BJ	Benin	HU	Hungary	PL	Poland
BR	Brazil	IE	Ireland	PT	Portugal
CA	Canada	IT	Italy	RO	Romania
CF	Central African Republic	JP	Japan	RU	Russian Federation
CG	Congo	KP	Democratic People's Republic of Korea	SD	Sudan
CH	Switzerland	KR	Republic of Korea	SE	Sweden
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovak Republic
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CS	Czechoslovakia	LU	Luxembourg	SU	Soviet Union
CZ	Czech Republic	MC	Monaco	TD	Chad
DE	Germany	MG	Madagascar	TG	Togo
DK	Denmark	ML	Mali	UA	Ukraine
ES	Spain			US	United States of America

NONBLOCKING MULTICAST FAST PACKET/
CIRCUIT SWITCHING NETWORKS

FIELD OF THE INVENTION

The present invention relates to switching networks,
5 and more particularly to nonblocking multicast fast
packet/circuit switching networks for routing information
from one end of the switching network to the other end.

BACKGROUND OF THE INVENTION

A new class of nonblocking multicast fast
10 packet/circuit switching networks has been invented. The
topology of these networks is based on the banyan
multistage interconnection network (See L.R. Goke and
G.J. Lipovski, "Banyan Networks for Partitioning
Multiprocessing Systems," First Annual Computer
15 Architecture, pp. 21-28, 1973) or its equivalent (See
C.L. Wu and T.Y. Feng, "On a Class of Multistage
Interconnection Networks," IEEE Trans. on Computers, vol
29, no. 8, pp. 694-702, 1980). These switching networks
route packets with proper multicast routing tags to their
20 destinations in a distributive and a nonblocking fashion.

The following discussion focuses on packet switching
although the invented switching architecture can also be
used for circuit switching.

For the introduction of ATM technology, the point-
25 to-multipoint services such as video distribution, video
conference, file transfer, distributed data processing,
and network management generate a lot of interest in
designing a multicast fast packet switch.

In general, there are five known approaches to
30 designing a multicast packet switch. The first is to use
a shared-medium switch. The second approach is to use
the store-and-forward nature of the packet switch, i.e.,
send the multicast packet one-by-one from the input port.

The third approach is to use a copy network and a routing network. The fourth approach is to use a multicast module at the output port and this multicast module is responsible for sending the multicast packet to the destinations. The fifth approach is to use a multicast banyan network or a multicast tree network.

Shared-Medium Approach

Due to the broadcast nature of the shared-medium, the multicast operation can be achieved without packet duplication. The general concerns of the shared-medium approach are the high bus speed and the complex media access protocol used to solve the bus contention. The bus speed requirement can be reduced by using a wider parallel bus. Nevertheless, it is not appropriate and cost effective to use the shared-medium switch, if the switch capacity is above 1 Gbits/s.

Store-and-Forward at the Input Port

In this approach, the multicast operation is achieved by sending the multicast packet one by one from the input port. The advantage of this approach is that a point-to-point switch can be used as a multicast switch; hence, the hardware cost for building a multicast switch is minimal. The disadvantages of this approach are the long delay due to the serial transfer of the multicast packet and serious congestion at the input port if the number of duplication is large.

Copy Network Plus Routing Network

In this approach, multicast packets are duplicated using a space-division copy network (See T.T. Lee, "Nonblocking Copy Network for Multicast Packet Switching," IEEE JSAC, vol. 6, no. 9, pp. 1455-1467, Dec. 1988.) (Also see J.S. Turner, "Design of a Broadcast

Packet Switching Network," IEEE INFOCOM, pp. 667-675, 1986, 1988). After the packet duplication, the point-to-point routing network routes the packets to the destinations. The difference between Turner's approach is that Lee's copy network is nonblocking and Turner's copy network is blocking. Since Turner's copy network is blocking. The switching elements in the copy network are buffered. Evidently, Lee's copy network is superior than Turner's copy network in several aspects: the non-buffered-banyan network, the nonblocking property, and constant latency time. For these reasons, the Lee copy network is used as the representative of this approach.

The first operation of the copy network is to duplicate the exact number of copies for each multicast packet. This requires the incoming packets to carry a copy index in the header. For the copy network to be nonblocking, the copy index of each multicast packet has to be translated into an address interval; this operation is performed on the incoming packets sequentially from the top to the bottom. This procedure can be implemented using a running adder network and an address interval encoder. After the multicast packets with proper address intervals have been generated, a concentration network is necessary to concentrate these packets to satisfy the nonblocking condition of the copy network. Since these packets' address intervals are monotonically increasing and they are concentrated, a banyan network can duplicate these packets without any blocking. After the copies of the multicast packets have been generated, a table is necessary to translate the header of each copy to the destination address so that the routing network can route the packet to the destined output port.

There are several disadvantages to using this approach. The first is the delay incurred for every packet (including unicast and multicast packets) passing through the copy network and the routing network, and the hardware complexity incurred by the copy network. Since the duplicated packets might use different input ports of the point-to-point switching network for routing at different times, these packets might be transmitted out-of-sequence due to different levels of input congestion at different times. It is noted that if the routing network is nonblocking, there is no out-of-sequence problem. For the same reason, if the positions of the duplicated packets at the input ports of the point-to-point routing network are changed very often due to connection addition or deletion, the translation table has to be updated very often also. This puts a large burden on the signaling processor. This situation becomes very severe for video distribution services since viewers will most likely change the channels very often. If the total number of duplication for each time exceeds the size of the switch, the copy network cannot handle this situation and the overflowed packets are simply dropped.

Multicast Module at the Output Port

In this approach, there are multiple modules at the output ports. All the multicast packets are relayed to these multicast modules first. Then the multicast modules send the multicast packet to the destined output ports through a proper multicast interconnection network. The number of multicast modules required depends on the amount of multicast traffic. The multicast knockout switch uses a similar approach (See K.Y. Eng, M.G.

Hiuchyj and Y.S. Yeh, "Multicast and Broadcast Services in a Knockout Packet Switch," IEEE INFOCOM, pp. 29-34, 1988.).

5 The disadvantage of this approach is that some of the output ports are allocated for the multicast modules. The result is that the switch may become a nonsymmetric switch if a symmetric switch is chosen as the point-to-point switching fabric. The other disadvantage is that there are two packet transfer protocols; one for the point-to-point network and the other for the multicast network.

Multicast Banyan or Tree Networks

15 The Multicast banyan network (see Figure 1) or tree network (see Figure 2) can duplicate and route the multicast packet simultaneously, which is very attractive.

20 In a multicast banyan network (See G. Nathan, P. Holdaway, and G. Anido, "A Multipath Multicast Switch Architecture," 1988.), the number of stages of 2×2 switching elements is $\log_2 N$ and the stages are numbered from 1 to $\log_2 N$. For example, in Figure 1, the number N of outputs is equal to 8 and therefore the number of stages is equal to 3. The format of the routing tag 101, which, along with data portion 102 makes up a packet 103, used for the multicast banyan network is different from the routing tag used for the point-to-point banyan network. The multicast routing tag 101 is formed using a series of 1's and 0's bits, and each bit is associated with one output port. Hence the size of the routing tag of a multicast packet is N . There are two registers 104 and 105, one for each output, holding control bits at each of the switching elements 106. The contents of the

control bits at each register is determined based on the topology of the banyan network. The operation of each switching element is to AND the routing tag 101 of the packet 103 and the control bits at each register (104 and 105). If the result of the AND operation is 1, a copy of the packet 103 is sent to the output. If the packet 103 is sent to both outputs, duplication of the packet has occurred. Although the multicast banyan network is cost-effective, it has the disadvantages of large registers in the switching element and serious internal blocking problem.

A tree network as shown in Figure 2 consists of two portions: a splitter network and a combiner network. The number of stages of 1 X 2 switching elements in the splitter is $\log_2 N$ and the number of stages of 2 X 1 switching elements in the combiner is also $\log_2 N$. Therefore, the total number of stages in a tree network is $2 \log_2 N$. $N=4$ in Figure 2. Although the tree network is internal nonblocking, it has the disadvantages of large registers in the switching element and high cost if the size of the switch is large.

SUMMARY OF THE INVENTION

It is found by modifying the topology of the multicast banyan network, a nonblocking multicast network can be constructed; hence, a family of nonblocking multicast switching networks based on the multicast banyan network is proposed. Compared with the blocking multicast banyan network, the new switching networks do not need long registers in the switching elements and they have no internal blocking problem.

Compared with the tree network, the new switching networks do not need long registers in the switching elements and the total hardware cost is lower.

5 Compared with the nonblocking multicast fast packet switch proposed by Lee, the invention does not separate the multicast operation into two separate phases, i.e., copy network and routing network. The invention routes and duplicates the packet at the same time in a distributive and a nonblocking fashion. The proposed
10 switch uses the input buffering approach and the locations of the multicast packets are fixed with respect to the input ports; hence, it does not suffer the disadvantage of the Lee approach, i.e., a frequent updating of the address translation tables at the input
15 ports.

A method is introduced below to create a nonblocking multicast banyan network. It is assumed that the input buffering approach is employed for the switching operation, i.e., the incoming packets are stored in the
20 input buffers. It is also assumed that a scheduling algorithm, such as the ring reservation scheme used (See B. Bingham and H. Bussey, "Reservation-Based Contention Resolution Mechanism for Batchier-Banyan Packet Switches," Electronic Letters, vol. 24, no. 13, pp. 772-773, June
25 1988.), is applied to the incoming packets at each slot time to reserve the output ports. The result is there is no output contention for the packet transfer through the switch at each slot time, i.e., no two packets destined to the same output(s) at the same time. To have a
30 consistent operation of the switching network, empty packets are generated at the input ports if no packets are ready to transmit at a slot time so that the total

number of packets at the switching network is always equal to the size of the switch.

The self-routing nonblocking multicast switching network according to the present invention, comprises:

5 input means for receiving a plurality of input messages, each of said messages potentially destined for a plurality of outputs of said switching network, each of said input messages containing a routing tag having a plurality of sections each section being composed of a
10 number of bits, said number of bits being different for each section within a routing tag, each of said sections corresponds to a level of a tree hierarchy related to the outputs of said switching network;

15 sorting means for sorting said input messages received by said input means by examining only one of said plurality of sections of each routing tag; and

20 routing means for routing the sorted messages output from said sorting means through to the outputs of said switching network based on the bits contained in said routing tags.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 shows an 8 X 8 multicast banyan network according to a well-known design;

25 Figure 2 shows a 4 X 4 tree network according to a well-known design;

Figure 3a shows a tree hierarchy of the routing field according to the present invention;

Figure 3b shows a general example of the routing tag format of the present invention;

30 Figure 3c shows an example of the routing tag format;

Figure 3d shows a second example of the routing tag format;

Figure 4 shows an 8 X 8 self routing nonblocking multicast banyan network;

5 Figure 5 shows an example of an 8 X 8 self routing nonblocking multicast banyan network;

Figure 6 shows a 16 X 16 nonblocking multicast switching network;

10 Figure 7 shows another 16 X 16 nonblocking multicast switching network;

Figure 8 shows still another 16 X 16 nonblocking multicast banyan network;

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS
Multicast Routing Field

15 The multicast routing field formats use the even and odd group concept associated with the levels of the switching network, and they are arranged using a tree hierarchy structure (see Figure 3a). The definition of a level in the proposed switching network will be explained later. At level 1, the even group (on the left-hand side of Figure 3a) consists of the output addresses whose modulo 2 results are 0; the odd group (on the right-hand side of Figure 3a) consists of the output addresses whose modulo 2 results are 1. The addresses at level 1 consist of 2 bits (301 in Fig. 3b) which are used for routing at level 1 of the switching network. There are four possible combinations of the 2-bit format: (1,1), (1,0), (0,1), and (0,0) which represent the destination addresses destined to both groups, even group, empty and odd group, respectively.

20
25
30

The addresses at level 2 consist of 4 bits 302 which are used for routing at level 2. The first 2-bit field

is associated with the even group at level 1 and the second 2-bit field is associated with the odd group at level 1. Examine the first 2-bit field 303. The subeven group within the even group at level 1 consists of the addresses whose modulo 4 results are 0 and the subodd group within the even group at level 1 consists of the addresses whose modulo 4 results are 2. Examine the second 2-bit field 304. The subeven group within the odd group at level 1 consists of the addresses whose modulo 4 results are 1 and the subodd group within the odd group at level 1 consists of the addresses whose modulo 4 results are 3.

In general, for a switching network with size N , the addresses at level m consist of 2^m bits, where $1 \leq m \leq \log_2 N$. The size of the multicast routing tag is $2N-2$.

The basic arrangement of the routing tag format is exemplified by concrete examples in Figures 3c and 3d. In Figure 3c, the outputs 0, 4, 6 and 1 are designated. In Figure 3d, the output 7 is designated.

Nonblocking Multicast Switching Network Operation

It can be observed that at stage 1 of the multicast banyan network there is no blocking if only one of the following three situations is allowed to occur at each switching element.

1. one packet which is destined to both groups and the other packet arriving at the switching element is an empty packet.

2. two packets where one packet is destined to one group (i.e., odd group) and the other is destined to the other group (i.e., even group).

3. one packet which is destined to only one group and the other packet is an empty packet.

In order to achieve the above objective, a sorting network is used to rearrange the pattern of the arriving packets. The sorting network sorts the packets using the 2-bit field 301 at level 1. Let the sorting network sort the packets into non-ascending order. After the sorting procedure, the sequence of the packets appearing at the outputs of the sorting network is packets destined for: both groups, even group, empty and odd group.

Using a well-known shuffle interconnection 401 to connect from the outputs of the sorting network 402 to the inputs of stage 1 of the banyan network, it is guaranteed that there is no blocking at stage 1 (see Figure 4). Although Figure 4 shows an 8 X 8 switching network, the approach can be applied to a switching network of any size. (A formal proof is given below in Appendix A.)

It has been shown that there is no blocking at level 1 of the network, where level 1 consists of one sorting network with size N and stage 1 of the banyan network.

The operation of each switching element at stage 1 of the banyan network is described as follows. The switching element referenced generally as 403 routes the packet to the upper link 404 if the 2-bit tag 301 is destined for the even group; it routes the packet to the lower link 405 if the 2-bit tag is destined for the odd group; it routes and copies the packet to both links if the 2-bit tag is destined for two groups. The empty packet is deleted if the other packet at the other input is destined to both groups; otherwise, the empty packet is sent to the next level. In summary, the 2-bit routing bits 301 at level 1 are used for sorting for the N-by-N sorting network and routing for stage 1 of the banyan

network. The switching element's logic is very simple, it only needs to check a 2-bit routing field.

5 After level 1, the packets have been divided into two groups according to the destination routing tags; the packets destined to the even group are routed to the upper subnetwork 501 and the packets destined to the odd group are located to the lower subnetwork 502 (see Figure 5). The level 2 portion 302 of the routing tag is used for routing at level 2 of the network which consists of 10 two sorting networks with size $N/2$ in parallel and stage 2 of the banyan network (see Figure 4). The upper subnetwork 501 in Figure 5 (or the lower subnetwork 502) consists of one sorting network 503 with size $N/2$ and the upper half (or the lower half) of stage 2 of the banyan network. 15

The upper subnetwork 501 with size $N/2$ uses the first 2 bits 303 at level 2 of the routing tag for routing, and the lower subnetwork 502 with size $N/2$ uses the second 2 bits 304 at level 2 of the routing tag for routing. The same routing procedure as in level 1 is 20 applied at each subnetwork.

This operation is repeated at every level until the last level. At the last level, the size of each subnetwork is 2. Hence, no sorting network is required 25 in this level. The last level of the network only consists of stage $\log_2 N$ of the banyan network.

The output ports of the switch check the routing tag of the arriving packet to determine whether it is an empty packet or not. If it is an empty packet, it will 30 be discarded. The logic to perform this operation is very simple, and only needs to check a 2-bit field.

For convenience, the proposed switching network has been named (N, n_1, n_2, p) network, where N is the size of the network, n_1 is the number of stages of sorting networks, n_2 is the number of stages of 2×2 switching elements, and p is the number of copies stacked together, where each copy is a $(N, n_1, n_2, 1)$ network. The proposed network can be described as the $(N, \log_2 N - 1, \log_2 N, 1)$ network.

The proposed switching architecture can be used for both packet switching and circuit switching. For circuit switching, since there is no output conflict and the switching fabric is multicast nonblocking, the result is a multicast nonblocking circuit switch. For packet switching, since the output conflict is an unavoidable situation, a scheduling algorithm is necessary to resolve the output contention. An example of the operation of an 8×8 nonblocking multicast switching network or an $(8, 2, 3, 1)$ network is provided in Figure 5.

Variations of the Nonblocking Multicast Switching Architectures

The banyan network has been shown to be topologically equivalent to many other multistage interconnection networks such as baseline, omega, flip, and shuffle networks [2] (See Wu Feng). Therefore, the proposed nonblocking multicast switching network can be constructed using any of the banyan network's equivalent topology.

The total number of stages of 2×2 switching elements is $\log_2 N$ and the total number of 2×2 switching elements is $N/2 \log_2 N$.

If the sorting network is implemented using the batcher sorting network (See K.E. Batcher, "Sorting

Networks and Their Applications," AFIPS, vol. 32, pp. 307-314, 1968), the total number of stages of sorting elements is

$$1/2[n(n+1) + (n-1)n + \dots + 2(3)], \text{ where } n = \log_2 N.$$

The total number of sorting elements is

$$1/4 N(n+1) + (n-1)n + \dots + 2(3)], \text{ where } n = \log_2 N.$$

It is found that the tree network topology can be combined with the nonblocking multicast $(N, n_1, n_2, 1)$ network to reduce the number of stages of the switching network. The principle is that one stage of the tree network is equivalent to one level of the multicast switching network; however, the number of copy of a switching network is doubled. The theory of the principle is explained as follows:

At level 1 of the network, a 1×2 demultiplexer is employed for each input line and these 1×2 demultiplexers can be substituted for the $N \times N$ sorting network and stage 1 of the banyan network. The outputs of the demultiplexer are connected to the inputs of two switching networks stacked in parallel (see Figure 6). The operation of the 1×2 demultiplexer is as follows.

For even number inputs, the multicast packet is separated into two packets: the packet with the even group goes to the first subnetwork of copy 1 and the packet with the odd group goes to the second subnetwork of copy 2.

For odd number inputs, the multicast packet is separated into two packets: the packet with the even group goes to the first subnetwork of copy 2 and the packet with the odd group goes to the second subnetwork of copy 1.

If there is no packet with the even group (or odd group), an empty packet is generated at the input. If there is no packet to send at a slot time, two empty packets are generated at the input. The principle is to make the number of packets at each copy of the network always equal to the size of the network.

With this operation, the packets are separated into two groups according to their destinations after stage 1. Each group is handled by two subnetworks, one in each copy. The operation at levels 2, 3 and 4 of each copy has been described above. After level 4, a combiner stage of 2 X 1 multiplexers is necessary to multiplex two outputs from two copies into one output. The resulting network becomes a $(N, \log_2 N-2, \log_2 N-1, 2)$ network. The operation of the 2 X 1 multiplexer is to discard the empty packet and send the data packet to the output port. The logic to perform this operation is very simple, which only needs to check a 2-bit field.

By the same token, if two stages of the tree network are provided, the sorting networks at levels 1 and 2 are not required (see Figure 7). After stages 1 and 2, the packets are separated into four disjoint groups according to their destination addresses. Each group is handled by four subnetworks, one in each copy. After level 4, two stages of 2 X 1 multiplexers is necessary to multiplex four outputs from four copies into one output. The network becomes a $(N, \log_2 N-3, \log_2 N-2, 4)$ network.

In general, an (N, n_1, n_2, p) multicast switching network is nonblocking if

- o $n_1 + \log_2 p = \log_2 N-1$ and
- o $n_2 + \log_2 p = \log_2 N$.

Appendix A

Theorem 1: The multistage sorting-banyan network is point-to-multipoint nonblocking.

Proof: Let the size of the switch be N . The
5 outputs consists of two groups, say G_1 and G_2 (see Figure
8). G_1 consists of the even addresses and G_2 consists of
the odd addresses. To make the total number of packets
presented to the sorting network always equal to N , the
empty packets are included in the sorting procedure. The
10 sorting network sorts the packets according to the
following methods. If the packet is destined to two
groups simultaneously, this packet is sorted to the top
of the output. After this, the packets destined to G_1
are put under the packets with two groups. Then the
15 empty packets are put under the packets destined to G_1 .
Those packets destined to G_2 are put at the end of the
outputs. Using a shuffle interconnection from the
sorting network to the inputs of stage 1 of the banyan
network, the packets are arranged in such a way that the
20 packets destined to the same groups do not appear at the
same switching element.

Let the number of packets destined to both G_1 and G_2
be N_0 , the number of packets destined to G_1 by N_1 , and the
number of packets destined to G_2 be N_2 .

25 Focusing on the packets destined to both groups and
the packets destined to G_1 . The sum of N_0 and N_1 is less
than or equal to $N/2$. If the sum is less than $N/2$, the
empty packets are used to pad the remaining positions.
The number of empty packets required is $N/2 - N_0 - N_1$.
30 According to the sorting procedure and the shuffle
interconnection, these $N/2$ packets can be put at the

first input of the switching elements at stage 1 sequentially from the top to the bottom.

5 The number of empty packets left is $N/2 - N_2$. If $N/2 - N_2$ is always greater than or equal to N_0 , then the packet destined to G_2 cannot appear at the same switching element with the packet destined to both G_1 and G_2 since empty packets will occupy the second input of the switching elements where the first input has a packet destined to both groups. Since the sum of N_0 and N_2 is
10 always less than or equal to $N/2$, $N/2 - N_2$ is always greater than or equal to N_0 . It has been shown that with this configuration, there is no blocking at stage 1.

After stage 1, the packets have been divided into two groups according to their destinations: even and odd.

15 Focusing on the even group G_1 at the upper subnetwork, G_1 consists of two subgroups G_{11} and G_{12} , where G_{11} has the addresses whose modulo 4 results are zero and G_{12} has the addresses whose modulo 4 results are two. Applying the same operation on G_1 at the upper subnetwork
20 as before, no blocking will occur at stage 2 of the banyan network. It is clear that this operation can be applied to all the stages recursively and the result is the network is point-to-multipoint nonblocking.

25 The true spirit of the invention is not to be limited by the above-described embodiments, but only by the appended claims.

WHAT IS CLAIMED IS:

1. A self-routing nonblocking multicast switching network comprising:

input means for receiving a plurality of input messages, each of said messages potentially destined for a plurality of outputs of said switching network, each of said input messages containing a routing tag having a plurality of sections each section being composed of a number of bits, said number of bits being different for each section within a routing tag, each of said sections corresponds to a level of a tree hierarchy related to the outputs of said switching network;

sorting means for sorting said input messages received by said input means by examining only one of said plurality of sections of each routing tag; and

routing means for routing the sorted messages output from said sorting means through to the outputs of said switching network based on the bits contained in said routing tags.

2. A network according to claim 1 wherein said sorting means includes first sorting means for sorting said input messages into the following order:

(a) input messages destined for both even-numbered and odd-numbered addresses

(b) input messages destined for only even-numbered addresses

(c) input messages destined for no addresses

(d) input messages destined for only odd addresses.

3. A network according to claim 1 wherein said routing means includes a plurality of switches.

4. A network according to claim 3 wherein each of said switches has two inputs and two outputs.

5. A network according to claim 4 wherein each of said switches includes:

odd-numbered addresses is received at one input and an input message destined for no addresses is received at the other input; and

second outputting means for outputting a message destined for no addresses on one output and a message destined for either all odd or all even numbered addresses on the other output when a message destined for no address is received as one input and a message destined for either all odd or all even numbered addresses is received as the other input.

6. A network according to claim 3 in which each of said switches performs its routing function by examining only two bits.

7. A network according to claim 1 in which each of said routing tags has a first section composed of two bits capable of forming four binary combinations, one combination indicating that a message is to go to only even-numbered outputs of said switching network, another combination indicating that a message is to go to only odd-numbered outputs, another combination indicating that a message is to go to both even- and odd-numbered outputs, and another combination indicating that a message is to go to no outputs.




Figure 1: 8 x 8 Multicast Banyan Network

Prior Art




Figure 2: 4 X 4 Tree Network

Page APT

Fig. 3a




Fig. 3b

Routing Tag Format:



Fig. 3c

Example: (0, 4, 6, 1)



Fig. 3d



X: don't care

Figure 3: Tree Hierarchy of the Routing Field




Figure 4: 8 X 8 Self-Routing Nonblocking Multicast Banyan Network [(8,2,3,1) Network]




Figure 6: Example of an 8 X 8 Self-Routing Nonblocking Multicast Banyan Network [(0,2,3,1) Network]




Figure 6: 16 X 16 Nonblocking Multicast Switching Network [(16,2,3,2) Network]




Figure 7: 16 X 16 Nonblocking Multicast Switching Network ((16,1,2,4) Network)




Figure 8: 16 X 16 Nonblocking Multicast Banyan Network

INTERNATIONAL SEARCH REPORT

International Application No.
PCT/US92/07978**A. CLASSIFICATION OF SUBJECT MATTER**

IPC(5) :H04L 12/56 H04J 3/24

US CL :370/60, 370/94.100, 370/94.300, 340/825.020

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 370/58.100, 370/60.1, 370/54, 370/68

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
<u>X</u> Y	US,A, 4,910,730 (DAY, JR. ET AL) 20 MARCH 1990 Columns (1-10), figures (1-20)	<u>1-3 and 6-7</u> 4 and 5
A	US,A, 5,018,129 (NETRAVALI ET AL) 21 MAY 1991 Columns (1-2), figure (4)	1
A	US,A, 4,701,906 (RANSOM ET AL) 20 OCTOBER 1987 Figure (2)	
A	US,A, 4,991,168 (RICHARDS) 05 FEBRUARY 1991 Columns (1-10), figures (1-5)	1-7

☐ Further documents are listed in the continuation of Box C.
 ☐ See patent family annex.

* Special categories of cited documents:	*T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be part of particular relevance	*X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	*Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*Z* document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means	
P document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search 09 NOVEMBER 1992	Date of mailing of the international search report 04 JAN 1993
Name and mailing address of the ISA/ Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231	Authorized officer DANG TON <i>Renewed for</i>
Facsimile No. NOT APPLICABLE	Telephone No. (703) 305-4739